

1 IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

2
3
4 APPLICATION FOR UNITED STATES PATENT

5
6 FOR

7
8 **SWITCH-BASED ACCELERATION OF COMPUTER DATA STORAGE**
9 **EMPLOYING AGGREGATIONS OF DISK ARRAYS**
10

11 Inventors: George Ericson; Robert Solomon; Jeffrey Brown; John Haynes
12
13
14
15
16
17
18
19

20 Attorneys:

21 Joel Wall, Esq.
22 P.O. Box 169
23 Hopkinton, MA 01748
24 508-435-4432
25

**SWITCH-BASED ACCELERATION OF COMPUTER DATA STORAGE
EMPLOYING AGGREGATIONS OF DISK ARRAYS**

FIELD OF THE INVENTION

The present invention relates generally to computer data storage systems and, more particularly, relates to acceleration of computer data storage utilizing fibrechannel switches, disk drive aggregators, and arrays of disk drives.

BACKGROUND OF THE INVENTION

Computer systems are pervasive in our society and virtually all human activity is now influenced at least to some extent by existence and usage of these systems. The faster and more efficient these systems are, the better for all concerned. Certain computer systems developing within the technological area known as fibrechannel or fibrechannel networks do offer faster and more efficient operation, not only because of their optically-communicative capability but for other reasons as well. One of the configurations in fibrechannel networks employs multiple disk drive arrays for data storage managed by an aggregator (essentially another array but with much higher intelligence than ordinary arrays and which organizes ordinary arrays into "aggregations") in combination with a fibrechannel switch (another intelligent device which performs a complex switching function under control of at least the aggregator). Typical inter-relationships of computer, aggregator, switch, and storage within fibrechannel networking have been established.

Referring to Fig. 1, there is depicted one typical fibrechannel computer system arrangement. Computer hosts 101, 102, and 103 communicate through fibrechannel switch or hub 104, sometimes known as a "fabric". The term "fabric" suggests densely-packed multiple conductors, since internal fibrechannel switch connections can be very dense. The irregularly-shaped "cloud" symbol representing the switch implies an active or changeable entity which is capable of being used or controlled. Front end fabric 104 connects to aggregator 105 (typically a RAID system, standing for Redundant Array of Independent/Inexpensive Disks) which, in turn, connects to back end fabric 106 (another fibrechannel switch or hub) to which are connected multiple disk drives 107, 108, 109, and 110. A major goal of this system is efficient movement of data or computer information from hosts to disk drive storage and vice-versa. If all computers communicate directly with all disk drives (and there can be many more than three hosts and four drives, those numbers being represented here only for purposes of clarity of illustration) then highly complex and inefficient operation with multiple hosts competing for the same storage space on the same disk drives, etc. can result. Thus, an aggregator is used to allow communication by computers with drives, but only through the aggregator to improve operation. The aggregator is a highly intelligent and complex device which appears to computers such as hosts 101, 102, and 103 to be a number of disk drives. The aggregator further appears to the computers to be the only disk drives in the system since it "hides" disk drives 107 - 110 connected to the back end fabric. This reduces complexity for computer hosts to a great extent. Further, this introduces a degree of security since all commands relative to data stored on disk drives from hosts must pass through, and thus be "approved" by, the aggregator. Any illegitimate command or

1 operation may be stopped by the aggregator before it does damage. But, unfortunately,
2 the aggregator can become a bottleneck in this configuration between computers and disk
3 drives under certain high-traffic or busy or other conditions. Thus, the aggregating
4 device can introduce "latency" or time delay into system operation and contribute to the
5 very inefficiencies in system operation that it was designed to reduce or eliminate. Under
6 certain circumstances, this can be a serious problem.

7 However, if the back end drives were directly accessible via the front end fabric,
8 the aggregation "bottleneck" would be removed and certain reductions in these latencies
9 might be achieved. In Fig. 2, Host computers 201, 202, and 203 are shown connected to
10 front end fabric - fibrechannel switch 204 to which are also connected aggregator 208 and
11 disk drives 205, 206, and 207. It is to be understood that the number of hosts and drives
12 are not limited to the specific number shown and that many more, or fewer, hosts and
13 drives are intended to be represented by this diagram. In operation, any one or more of
14 the hosts first sends data requests to the aggregator which then enables the disk drives
15 and alerts them that these requests are coming directly to any one or more of them. Then
16 hosts send multiple requests addressed to the disk drives through the switch *directly* to
17 these different drives, accessing these drives in parallel and receiving *directly* back
18 multiple data streams in parallel through the switch, which reduces the latency factor by
19 eliminating at least one "hop" through the aggregator. However this configuration re-
20 introduces the security issue, because these drives, not being "protected" by the
21 aggregator, are more exposed to illegitimate commands in this configuration. Thus, disk
22 drives and computers in this configuration have to contain added intelligence to deal with

1 these security issues and the task of adding this intelligence creates a more complicated
2 and less desirable environment.

3 Referring next to the subject of fibrechannel protocols as further useful
4 background information, a book entitled "Fibre Channel Volume 1 The Basics" by Gary
5 R. Stephens and Jan V. Dedek, published by Ancot Corp, Menlo Park, California, first
6 edition June, 1995, is incorporated by reference herein. Within the computer industry
7 there are highly competitive companies which specialize in design, development and
8 manufacture of these switches, aggregators, memory arrays, and other fibrechannel-
9 related components. If their respective designs are to be employed in the same system, or
10 if multiple systems employing their various designs are networked together, these designs
11 have to mesh together properly for users to derive any benefit from them. This is
12 accomplished by having these companies agree to certain standards sometimes
13 generically known as the "ANSI Fibre Channel Standard". These standards are complex
14 and are negotiated into existence by the very companies that are responsible for creating
15 these fibrechannel-related components. One of the agreedupon products of these
16 negotiations is what is sometimes called the "protocol stack"-- five network levels of
17 fibrechannel. (In computer networks, information or data sent between network devices
18 is conducted on a physical level normally by electrons or photons over copper wires or
19 fibre-optic paths respectively, and/or by telecommunication paths, and, at the same time,
20 is also *virtually* conducted on multiple other network levels above the physical level.)
21 Referring to Fig 3 A, five levels: FC-0, FC-1, FC-2, FC-3 and FC-4 are shown,
22 corresponding to: physical, encode/decode (8B/10B), Framing Protocol, Common
23 Services for Ports, and Mapping respectively. {Sometimes, another sixth layer, Upper

1 Layer Protocol, is referred to, and is shown. } Briefly, the FC-0 functional level relates to
2 physical connection of nodes, either optical or electrical – the nuts and bolts of
3 connection. The FC-1 functional level relates to how information is transmitted between
4 fibrechannel input/output ports, i.e. how lasers and electrical drivers/receivers deal with a
5 bit stream moving into and out from a fiber. The FC-2 functional level deals with
6 transferring information and is concerned with its content, proper arrival of content or
7 detection of missing information or information errors; this level thus defines frame fields
8 including frame header field layout and is utilized in embodiments of the present
9 invention. The FC-3 functional level deals with common services that can be shared
10 among ports. And, the FC-4 functional level handles mapping of existing non-
11 fibrechannel I/O interfaces for use on fibrechannel by using fibrechannel tools.

12 The foregoing latency problem of the prior art is addressed and relieved, without
13 reducing security, by the welcome arrival of the present invention which operates not
14 only within parameters of the ANSI Fibre Channel Standard, but, as suggested, makes
15 novel use of fibrechannel level FC-2, as described hereinbelow.

17 SUMMARY OF THE INVENTION

18 The present invention in a broad aspect relates to a network-attached storage
19 computer system having disk drives and an aggregator attached to the network. Direct or
20 indirect data transfer between computer and disk drives is determined by its impact on
21 overall performance of the system. If determined that indirect transfer would increase
22 overall system performance compared to direct transfer, data is sent between computer
23 and disk drives through the network *and* through the aggregator. If determined that direct

1 transfer would increase overall system performance compared to indirect transfer, data is
2 sent between computer and disk drives through the network but *not* through the
3 aggregator.

4 The present invention in another aspect relates to a computer data storage system
5 wherein data is grouped in frames. There are disk drives or the like for storing and
6 retrieving data and an aggregator or data storage manager for managing operation of the
7 disk drives. Each of the frames includes a frame header which designates parameters
8 associated with data in its frame. One of the parameters is destination ID (identity of the
9 destination). There is a controllable switch connected between computer, disk drives,
10 and aggregator for selecting certain frames and flowing data in those selected frames
11 directly between computer and disk drives. The aggregator is destination ID in the
12 selected frames, but transfer of data in the selected frames between computer and disk
13 drives is direct and not through the aggregator. Thus, latency issues can be reduced or
14 eliminated because of direct flow of data between computer and disk drives, while at the
15 same time not reducing security since the destination ID for data in the selected frames
16 remains the aggregator.

17 In a further feature of the present invention, the switch includes switch control
18 logic under command of the aggregator to select those frames to be transferred directly.
19 The switch control logic includes a frame header field selector such as a frame header
20 mask, an input frame header buffer, and a map table.

21 In yet another aspect, the present invention is incorporated in a computer system
22 including both disk drives or the like for storing and retrieving data grouped in frames
23 and an aggregator normally in the path of the frames flowing between computer and disk

1 drives. The present invention employs computer logic and relates to enhancing transfer
2 of data between computer and disk drives. This involves establishing a frame header
3 field selector such as a mask containing only relevant information. The frames including
4 their respective headers are received from the computer, and headers and mask are
5 compared to obtain "distilled" frame headers. A map table is established which contains
6 sets of frame header fields corresponding to input addresses (proxy destination IDs) of
7 the disk drives. The map table is searched to find matches between distilled frame
8 headers and sets of frame header fields. For each match, a proxy destination ID is
9 substituted in place of the original destination ID in headers of each of the corresponding
10 frames which are then forwarded directly to the disk drives and not via the aggregator.
11 This data transfer enhancement operation is not perceptible by the host computer.

12 And in yet another aspect, the present invention relates to a computer program
13 product for use in a computer system employing network-attached storage having both
14 disk drives and a disk drive aggregator attached to the network. There is included a
15 computer usable medium having computer readable program code thereon for enhancing
16 the transfer of data between the computer and the disk drives.

17 It is thus advantageous to use the present invention to reduce latency issues
18 without negatively impacting data security in a network-attached-storage-based computer
19 system.

20 It is a general object of the present invention to provide increased overall system
21 performance in a computer system.

1 If is a further object of the present invention to provide improved performance in
2 storage and retrieval of data in a computer system, including a network-attached-storage
3 computer system.

4 It is an additional object of the present invention to provide apparatus and
5 methodology for allowing direct data flow between a computer system's host computers
6 and disk drives under certain conditions while maintaining data security.

7 It is a still further object of the present invention to provide an improved
8 fibrechannel-based computer system employing multiple disk drives wherein latency
9 normally introduced by bottleneck-operation of an aggregator is reduced or eliminated.

10 Other objects and advantages will be understood after referring to the detailed
11 description of the preferred embodiments and to the appended drawings wherein:

12 13 **BRIEF DESCRIPTION OF THE DRAWINGS**

14 Fig. 1 is a diagram of a prior art arrangement of host computers, fibrechannel
15 switches, aggregator, and disk drives;

16 Fig. 2 is a diagram of another prior art arrangement of host computers,
17 fibrechannel switch, aggregator, and drives;

18 Fig. 3A depicts the fibrechannel levels with their respective designations;

19 Fig. 3B depicts a layout of a typical frame in accordance with fibrechannel level
20 FC-2;

21 Fig. 3C depicts a layout of at least a portion of the fibrechannel level FC-2 frame
22 header of Fig. 3B;

Fig. 4 is a schematic diagram of a write command processed in accordance with an embodiment of the present invention;

Fig. 5 is a flowchart depicting the algorithmic process performed by the embodiment of the present invention of Fig. 4;

Fig. 6 depicts a map table listing certain sets of frame header fields associated with their respective proxy destination IDs;

Fig. 7 is a schematic diagram of the switch of Fig. 4 showing its switch control logic including the map table of Fig. 6 under command of the aggregator of Fig. 4; and,

Fig. 8 is a schematic diagram of a read command processed in accordance with an embodiment of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 3B

Referring to Fig. 3B, a typical fibrechannel frame in accordance with level FC-2 of the fibrechannel protocol stack is shown. This frame contains multiple fields which normally contain predetermined ranges of bytes. Typically the opening “idles” field has six transmission words or 24 bytes; “start of frame” (SOF) field has one transmission word or four bytes; the “header” field has six transmission words or 24 bytes and is of primary interest because it relates to the present invention and contains information about its frame’s contents or purpose, somewhat analogous to ID/origin/destination information displayed on a human-transit bus above the bus driver’s windshield; the “optional headers and payload” field can be the largest field ranging from zero to 528 transmission words or 2112 bytes; “cyclic redundancy check” (CRC) and “end of frame” (EOF) fields

are each one transmission word of four bytes each; and closing “idles” field is again six transmission words or 24 bytes.

Figure 3C

Referring to Fig. 3C, the frame header of Fig. 2 is expanded to show the five words of interest in this header, the number of bytes per header (four) and various fields associated with particular bytes. The header actually contains six words. Starting at the left-hand side, R_CTL is “routing control” represented by byte zero. D_ID is “destination identity” represented by bytes 1-3. Byte 4 is not used. S_ID is “source identity” and is designated by bytes 5-7. Byte 8 is “type”. F_CTL is “frame control” and is designated by bytes 9, A, B. SEQ_ID is “sequence identity” and is byte C. DF_CTL is “data field control” and is byte D. SEQ_CNT is “sequence count” and is designated by bytes E and F. OX_ID is “originator’s identity” and is designated by bytes 10-11. Finally, RX_ID is “receiver’s identity” and is designated by bytes 12-13. The D_ID, S_ID, and OX_ID fields in this header are utilized by embodiments of the present invention. Other fields may also be used and are selected by the frame header field selector such as a mask or its equivalent. These specific fields and these other fields shall be described more fully hereinbelow.

Figure 4

Referring to Fig. 4, a schematic block diagram of a computer system operating in accordance with an illustrative embodiment of the present invention is presented. Host computer 401 is labeled “I” for “initiator”; fibrechannel switch 402 is identified and shown as a box rather than a cloud only for purposes of convenience of illustration; aggregator 403 is labeled “A”; memory array or disk drive group 404 is labeled “Z”, and

1 memory array or disk drive group 405 is labeled “Y”. These system components are
2 shown as being interconnected by communication paths which are syntactically identified
3 for purposes of ease of description. The syntax is as follows:

4 Command N(S_ID, D_ID, OX_ID)(Proxy x_ID, x)

5 where “Command” is either a “read” from or “write” to memory with respect to host 401,
6 “N” is a number representing communication count in this particular group or exchange
7 of communications, “S” represents source of the command, “D” represents destination of
8 the command, “OX” represents originator’s exchange, “Proxy x” represents a substituted
9 source or destination command where “x” represents either substituted source S or
10 substituted destination D, and “ID” stands for the respective component’s identity in each
11 case.

12 Write Command Operation

13 In operation, host computer 401 issues a write request or command with syntax
14 label W1(I,A,1) which stands for: W = Write command; 1 = this is the first
15 communication in this particular exchange of communications designated “1” and
16 pertaining to this Write command; (I = S_ID; A = D_ID; 1 = OX_ID). Switch 402
17 receives W1(I, A, 1) at one of its input ports (not shown in this Fig.) and transfers it
18 through to aggregator 403 via communication path syntactically labeled W2(I,A,1) which
19 is essentially a continuation of W1(I,A,1). In response to W2(I,A,1) aggregator 403
20 issues two commands: a mapping command and a proxy write command. Map command
21 M3(I,A,1)(Z,D) is forwarded to switch 402 over command input port 406 where it
22 commands a mapping function within the switch (to be discussed below). Proxy write

1 command PW4(A,Z,1)(I,S) is also forwarded to switch 402 where it is switched to disk
2 drive Z via the line syntactically labeled PW5(A,Z,1)(I,S).

3 Before proceeding further with the syntactical narrative of Fig. 4, consider what
4 has been thus far represented: A write request is issued by the host computer – it wishes
5 to write data into a disk drive. It sends the request to the aggregator (via the switch)
6 which is in charge of efficiently managing disk drive resources available to the computer.
7 Upon receipt of this write request, the aggregator not only “decides” that this request
8 should be directed to disk drive Z rather than disk drive Y, but also “decides” that it
9 would be in the system’s best interests if future communications from Host 401 with
10 respect to this particular write request *not* proceed through the aggregator. Therefore, the
11 aggregator issues a map command to the switch to generate a frame header mask and a
12 map table (both shown in Fig. 7 and to be discussed below) for purposes of diverting
13 certain future communications from initiator 401 to substitute or proxy destination disk
14 drive Z (also to be discussed in detail below). The aggregator further issues a Proxy
15 write command PW4(A,Z,1)(I,S) to communicate to disk drive Z a proxy or substitute
16 source ID, namely that of initiator “I” rather than its own ID. Accordingly, at this point
17 in the communication proceedings for this write command, (1) the switch has been
18 commanded to forward future communications from I for this write command having
19 “A” as destination, from I *directly* to disk drive Z without proceeding through A, and (2)
20 disk drive Z has been “advised” or commanded to substitute initiator “I” as the source for
21 write requests which arrive at Z actually by way of A.

22 Returning, now, to the syntactical description of Fig. 4, disk drive Z responsive to
23 proxy command PW5(A,Z,1)(I,S) issues a transfer ready signal or command identified as

1 X6(A,I,1) which is interpreted as follows: "X" means transfer; "6" is the number
2 corresponding to this sixth communication step in the same originator's exchange
3 process; (source ID is A; destination ID is I; and this is still the same Write request
4 identified by originator's exchange "1"). This transfer ready command goes through
5 switch 402 and via the communicative link syntactically identified as X7(A,I,1) is
6 transferred to its destination, Initiator I, "advising" initiator I that conditions are ready for
7 data to be written. Initiator I, responsive to this transfer ready command, sends data to
8 switch 402 via the communicative link syntactically identified as D8(I,A,1), which stands
9 for: data, the eighth step count in this process, (source is I, *destination is A*, and same
10 write request #1) respectively. The destination was italicized in the prior sentence to
11 emphasize that A shall not receive this data despite destination "A" being designated in
12 the syntax because of map conditions earlier established in the switch (which will be
13 discussed below). In response to the map operation in the switch to be described below,
14 the switch detects a match involving disk drive Z and certain frame header elements,
15 substitutes disk drive Z for aggregator A 403 in the header's D_ID field, and sends this
16 data *directly* to disk drive Z via communicative link identified by D9(I,Z,1). Disk drive Z
17 then completes the proxy write command by issuing back to aggregator A a status signal
18 identified as S10(Z,A,1) where Z is source, A is destination, etc. This status signal goes
19 through the switch and via the link identified by S11(Z,A,1) from the switch to A where
20 it provides a status report to A that the write was successfully completed. At this point, A
21 can dismantle or neutralize the map table it established in the switch, whereupon it issues
22 an "unmap" command UM12(I,A,1) to the switch and the map table is dismantled
23 whereby its operational effect is neutralized.

1 In other words, in the last paragraph and the operation of Fig. 4 thus far, it should
2 be understood that the transfer ready signal advised the computer that conditions are
3 ready for data to be transferred, and that the computer sent the data addressed to the
4 *aggregator A*, via D8(I,A,1). Because of the switch's intelligence it "decided" to not
5 allow this write command to flow to the aggregator, and diverts it by way of the map
6 table directly to disk drive Z on link identified as D9(I,Z,1). After the status report is
7 made from the disk drive via the switch to the aggregator, the aggregator decides to
8 dismantle the map table which it no longer needs for this particular write command, via
9 unmap command UM12(I,A,1).

10 It can be seen that *all* destination ID's for host-originated commands are the
11 aggregator A, as in (note the bold emphasis) W1(I,**A**,1) and D8(I,**A**,1), and, *all* source
12 ID's for host-received commands are aggregator A, as in (again note bold emphasis)
13 X7(**A**,I,1) and S14(**A**,I,1). These results obtain even though the transfer ready command
14 originates on disk drive Z and not A and even though the data command causes data to be
15 written into disk drive Z and not A. Therefore, *from the host's viewpoint, nothing has*
16 *changed with respect to the switch or the aggregator with which it is connected!* This
17 entire proxy operation in redirecting commands directly to specific disk drive
18 components is not perceptible by the host.

19 All commands and signals except for possibly the map/unmap and proxy
20 commands are standard fibrechannel level FC-2 compatible commands. Accordingly,
21 since the syntax of all information transmitted in Fig.4 (except for possibly map/unmap
22 and proxy commands) is solely reflective of fields in the frame header, all such
23 information with possible exceptions noted can fit within the frame header. This

embodiment of the present invention takes advantage of the potential afforded by the frame header to insert proxy commands into the header to achieve this important result in latency reduction. Typically, map/unmap commands can be fabricated at the SCSI (small computer system interface) level, or could be at the microcode level in assembler language; alternatively, they can also be FC-2 compatible. In any case, they would be compatible with the fibrechannel switch internals. Furthermore, the proxy commands are proprietary SCSI upper level protocol commands which map to SCSI lower level protocol commands, which are, in turn, mapped onto communications media dependent protocols such as the FC-2 protocol layer.

All write command activity shown in Fig. 4 was executed with respect to disk drive Z because of a decision made by the aggregator to use disk drive Z. However, it could have selected disk Y or, alternatively, could have decided to alternate host-initiated write and/or read commands between disk drives Z and Y (and any other disk drives in the system not shown in this Fig.) and permit the system to process these commands virtually in parallel. In the latter alternative case the map operation (to be discussed below) could be extended to include sequence count and sequence ID information from the frame header where the map operation would be more complex. Such syntax would take the following form:

Map(S_ID, D_ID, OX_ID, SEQ_ID, SEQ_CNT)[(Proxy x_ID), x]

where all items have been previously identified.

Figure 8 – Read Command Operation

A read command example would be similar to the write command example shown, but a map command is not needed because the switch does not need to be

1 involved in the proxy operation for read commands. The transfer ready command is also
2 not needed in a read example because the relationship between (initiator) host computer
3 and (target) disk drive is that the latter drives the sequencing. The host initiator is
4 assumed to be ready, since it initiated the read request. So after the read request is issued,
5 whenever the (target) disk drive is ready to send data, it does. Referring to Fig. 8, there is
6 presented the same schematic block diagram of a computer system as used in Fig. 4, and
7 employing the same syntax, but where disk drive Y, drive 405, is used in this read
8 example instead of reusing disk drive Z, drive 404, for purposes of clarity of presentation.
9 As in the write example, Initiator 401 forwards a read command syntactically identified
10 as R1(I,A,2) with the usual meaning: R = a read request, 1 = the first command in this
11 series, (I = Initiator is the source of the command, A = Aggregator is the destination of
12 the command, and 2 = the second originator's exchange). This read request flows
13 through the switch and via a link identified syntactically as R2(I,A,2) is received by
14 aggregator 403. The aggregator responds by turning the read command into a proxy read
15 request: PR3(A,Y,2)(I,S) which travels through the switch to disk drive Y via the link
16 identified by PR 4(A,Y,2)(I,S). In this proxy command, PR = proxy read; 4 = the fourth
17 command of this originator's exchange; (A = actual source aggregator; Y = actual
18 destination disk drive Y; 2 = second originator's exchange)(I = Initiator as the proxy; and
19 S = "Source", meaning that the aggregator is identifying or substituting the proxy
20 Initiator in the role of source to disk drive Y). Disk drive Y retrieves the data from
21 within its storage and delivers it to the switch over the link identified by syntax:
22 D5(A,I,2), where, as reflected by "A" being in the source position of the syntax, disk
23 drive Y is responding in the role of aggregator as source of the data command thus

1 conforming itself to the appropriate component that would have just received a command
2 from the Initiator, namely, the aggregator. The switch forwards the data to I via the link
3 identified by D6(A,I,2). (Note that "A" is still designated the source although "Y" is the
4 actual source of this data.) Drive Y issues a status confirmation via link identified by
5 S7(Y,A,2) through the switch and via the link identified by S8(Y,A,2) to the aggregator.
6 The aggregator issues a status confirmation to the Initiator via the link identified by
7 S9(A,I,2) through the switch and via the link identified by S10(A,I,2) to the initiator. The
8 major difference(s) between this read example and the prior write example is that the
9 switch does not need to redirect any commands in the read example. Therefore the map
10 command is not needed (and thus the unmap command is not needed).

11 Focusing on the proxy read command PR3(A,Y,2)(I,S) or its continuation
12 PR4(A,Y,2)(I,S) and the responsive data command D5(A,I,2) or its continuation
13 D6(A,I,2), the significance of the functionality expressed by the syntax should not be
14 overlooked. These commands mean that the *aggregator* (which would have otherwise
15 been the source with respect to disk drive Y in an ordinary read command) is sending this
16 *read* command to the *disk drive Y* (which would have otherwise been the destination with
17 respect to the aggregator in an ordinary read command) and is herewith syntactically
18 identified within the proxy command ***PR3(A,Y,2)(I,S)*** in bold italics. But, this command
19 is a proxy command and is thus sending *additional* information identified in the proxy
20 portion of the command "***PR3(A,Y,2)(I,S)***", again in bold italics. The additional
21 information is a request to *substitute the initiator as the source of this command*. This
22 complete proxy command is received by Disk drive Y which *cooperates* and makes that
23 substitution. The result of this cooperation is observable in the responsive data command

output by disk drive Y. Disk drive Y sends the data *not* back to A, the actual source of the command, but to I, the proxied or substituted source of the command, *and identifies itself in the data command as aggregator A*. Examining the syntax for the data command one observes that *the aggregator is in the position of the source, and the initiator is in the position of the destination*. The net effect of these substitutions is that: (1) an additional hop through the aggregator is avoided when data is forwarded from a disk drive responsive to a read request from the initiator – data goes directly to the initiator (albeit through the switch or network) from the drive rather than (again via the switch) back through the aggregator from which the command actually was received, and (2) the host initiator is not impacted since the commands it sends and receives do not suggest anything other than what the host had expected – requests made by and directed from the initiator host to the aggregator as a destination, and data received by the host initiator from the aggregator as a source! Thus, “the host is not in the game”, using a colloquial expression to highlight the fact that other components or sub systems in this overall network attached storage system are cooperating or conspiring in a manner to improve throughput and improve other system performance characteristics without the host being aware of these changes, or substitutions, or proxies. And, as noted with the earlier-discussed write example, this entire proxy operation is not perceptible by the host.

Figure 5

Referring next to the flowchart of Fig. 5, it depicts the algorithmic process inherent in operation of switch 402 in the computer system of Fig. 4 for the write command example illustrated. An input frame header (including its complete frame for which the header is identifying information) from the host computer enters an input port

1 on the switch in block 501. It is to be understood that there can be multiple hosts each
2 sending write and/or read commands to this complex switch on its multiple input ports
3 essentially simultaneously, and the switch under control of its aggregator shall be able to
4 process all inputs appropriately. (This singular example of a write command is hereby
5 disclosed for purposes of enhancing clarity of presentation and understanding of
6 operation of the present invention. The corresponding flowchart for the read example
7 would be less complex than this because there is no map table operation associated with
8 the read command.) The algorithmic process moves to block 502 wherein a logical
9 “AND” is performed between the input frame header received and a frame header field
10 selector such as a frame header mask. The mask (more detail of which is presented
11 below in connection with Fig. 7) is a template allowing only information relevant to this
12 process. In other words, there are fields in the fibrechannel FC-2 frame header that may
13 contain information irrelevant to operation of the present invention, and they are filtered
14 out. The result of this logical “AND” step is information equal to or less than the mask
15 information, i.e. a subset of the mask information termed a “distilled frame header”.
16 The algorithmic process moves then to decision block 503 wherein the query is
17 presented: is there an entry (i.e.: a set of frame header fields) in the map table matching
18 the distilled frame header? This map table, earlier referred to in connection with
19 description of Fig. 4 and to be described in more detail in connection with Fig. 7
20 hereinbelow, is a dynamic table established in the switch. If the answer is “no”, the
21 process moves to block 506 which routes the entire frame associated with the distilled
22 frame header to the destination identified by the original D_ID value in the header, which
23 is a particular output port on the switch and the process is done. On the other hand, if the

1 answer is “yes”, then there is a match between the distilled frame header and a particular
2 frame header entry in the map table and the algorithmic process moves to block 504
3 where a substitution takes place. The particular D_ID value corresponding to the
4 matching entry (set of fields) in the table is substituted into the D_ID field in the input
5 frame header in place of the original value thus obtaining a “proxy” frame header. In
6 other words, the result of this operation is a changed input frame header associated with
7 the write command, this change taking place within the switch: the frame header was first
8 “distilled” whereby irrelevant header information for this process was removed, and then
9 the input frame header had its destination changed to this proxy destination. (In
10 summary, Proxy D_ID replaces the header field D_ID and Proxy S_ID replaces the
11 header field S_ID when they are used.) The algorithmic process moves next to block 505
12 where the complete input frame, for which the proxy frame header is its identification, is
13 routed to a *different* particular output port on the switch corresponding to the proxy frame
14 header’s D_ID value and the process is done. This means that data associated with this
15 write command will be sent to a destination different from that originally intended,
16 namely *directly* to a disk drive rather than the aggregator, and this change will take place
17 within the switch.

18 Figure 6

19 Referring next to Fig 6, map table 600 is presented. This is a table which exists
20 dynamically in fibrechannel switch 402. In other words, this table can be created, table
21 entries can be changed, and the table can be deleted by operation of hardware, firmware,
22 and/or software in the switch and controlled by the switch which is, in turn, controlled by
23 aggregator 403. More detail on this operation will be given with respect to Fig. 7

hereinbelow. The table shown has two main columns, the one on the left being sets of Frame Header Fields and the other on the right being Proxy D_ID's. In the Frame Header Fields column are sub columns for S_ID, D_ID, and OX_ID. More sub columns could be entered as may be needed, as for example, entering SEQ_ID and SEQ_CNT to accomodate the disk Z/disk Y alternating scheme described earlier where sequence identity and sequence count are additional fields needed to implement the alternating scheme. Furthermore, this particular table as depicted contains entries that are reflective of the write operation described in connection with Fig. 4 and Fig. 5. The first row, for example, shows a set of earlier discussed values for S_ID, D_ID and OX_ID being associated with a Proxy destination ID known as D_ID"a". Other sets of entries in the table may show different values for S_ID reflective of multiple computer hosts (up to "n" computer hosts) and different values for OX_ID reflective of multiple write commands per host (up to "M" number of originator's exchange write commands with computer host "1", up to "N" number of originator's exchange write commands with computer host "2", and up to "P" number of originator's exchange write commands with computer host "n"). Each of these sets of entries is associated with a particular destination ID as suggested by the column heading, e.g. a particular disk drive. Thus this column under Proxy D_ID's contains a range of destination IDs running; from D_ID"a" to D_ID"zzz" as shown. This nomenclature is intended to be suggestive of a large number of destination IDs and no particular fixed or limited number is intended to be represented hereby. Each of these destination IDs necessarily maps to a particular output port on the switch which connects to an appropriate disk drive. More discussion about this table will be provided in connection with Fig. 7.

Figure 7

Referring next to Fig. 7, switch 402 is shown with input ports 703, output ports 704 and 705, and containing, among other things, its switch control logic 700. Among other logic components not shown, switch control logic contains map table 600, frame header field selector or mask 701 and input frame header buffer 702. Map table 600 is the table described in Fig. 6. Header mask 701 and buffer 702 are dynamic constructs which can be implemented or dismantled, and are fashioned from computer hardware, firmware, and/or software within and/or related to the switch. Aggregator 403 commands control logic 700 by way of control port 406. Those skilled in the art and familiar with fibrechannel switches will appreciate how such control logic may be implemented from hardware, firmware and software.

In operation, if a write command is forwarded by a host to switch 402, it enters on one of input ports 703. This command is a full frame in fibrechannel protocol as shown in Fig. 3B. As it enters the switch, its header field as shown in Fig. 3C is loaded into header buffer 702, and if there is no prior-established proxy condition where there is no proxy match, the full frame is conducted to aggregator 403. Aggregator 403, as an overseer or manager for multiple disk drives, is aware of the state of activity and traffic regarding all of its disk drives, and makes a decision regarding if it would be in the computer system's overall best interest to (1) have this write command pass through itself to the disk array and then have such command's associated data also pass through itself on its way to the disk drive array, or (2) arrange for a proxy ID so that such data shall pass from the host computer into the switch and then written *directly* into the disk to avoid a bottleneck in the aggregator. Assuming the aggregator decides for a proxy ID, it

1 sends a map command via control port 406 to switch 402 which requests the switch to
2 fabricate input header mask 701. The aggregator signals to the switch in this command
3 precisely what elements should be put into the mask based on existing conditions and on
4 the requirement of handling a write command from the host. Additionally, the aggregator
5 also had commanded the switch control logic in this map command to fabricate map table
6 600 with specific frame header field sets of entries with their corresponding proxy
7 destinations based on existing conditions and on the requirement of handling a write
8 command. Thereupon, the aggregator commands the switch control logic to compare the
9 distilled header mask with the frame header field map table's sets of entries to seek a
10 match and to select the Proxy D_ID associated with that matched entry set as the new
11 switch destination ID for data to be sent by the computer and to be written into the disk
12 drive array. Accordingly, when data is sent by the host [syntax D8(I,A,1) in Fig. 4] to
13 aggregator 403 responsive to a transfer ready command from the specified disk drive
14 [syntax X6(A,I,1)] it first goes to the switch and then by prearranged proxy just described
15 goes directly to disk drives and not to the aggregator.

16 Those skilled in this art including those directly connected with design and
17 development of fibrechannel switches will recognize the fact that implementation of
18 illustrative embodiments of the present invention is within their skills and expertise and
19 will utilize appropriate hardware, firmware, and software to generate the logic to
20 accomplish these implementations. For example, a typical host computer which can be
21 used in connection with the present invention is any Intel, Sun Microsystems, Hewlett
22 Packard or other similar company's computer using a fibrechannel host bus adapter with
23 fabric support. A typical fibrechannel switch which might be used in connection with the

1 present invention and which can be fully implemented for Map/Unmap operations can be
2 obtained from companies such as Brocade, McData, Vixel, or Ancor, etc. Typical disk
3 drives which can be used in connection with the present invention are: any fibrechannel
4 disk drive modified to support proxy read and proxy write. The frame header buffer, map
5 table, frame header mask, and map/unmap commands would typically be implemented by
6 combination of software and hardware in or related to the switch. Aggregator 403 can
7 typically be derived from EMC CLARiiON model nos. 4700 which could have the
8 capabilities called for herein. Furthermore, the specific illustrative embodiments
9 presented are not the only ways of fabricating the present invention. For example, other
10 ways of utilizing a fibrechannel switch to accomplish the goals of the present invention
11 include use of hash table lookups for efficient access to a map table.

12 In a hash table lookup design, decision block 503 in Fig. 5 could be a hashing
13 algorithm. Hashing can take different forms. One form could be to mask off all bits in a
14 field except some bits; the unmasked bits could be either high, or low, or mid order bits.
15 The unmasked bits are used as an index, as an approximation to get to the general area of
16 the correct answer very quickly. An alternative hashing design would take all bits in the
17 field but would fold them together to obtain a smaller number of bits, again to accomplish
18 the very fast approximation objective. There are other hashing designs as well. In any of
19 these hashing designs one can avoid an exhaustive search, entry by entry, using the
20 distilled frame header and each entry set of fields in the map table, as earlier presented.
21 A hashing approach in connection with such a table would provide an output advising
22 whether or not any matches existed in the sampled subset of the table – and if not then the
23 search could move on quickly to the next sampling of entries.

1 The present embodiments are therefore to be considered in all respects as
2 illustrative and not restrictive. For example, the invention need not use a fibrechannel
3 switch; any functionality that is the equivalent of such a switch, such as Infiniband could
4 be utilized with the present invention. The scope of the invention is indicated, therefore,
5 by the appended claims rather than by the foregoing description, and all changes which
6 come within the meaning and range of equivalency of the claims are therefore intended to
7 be embraced therein.